



## The Who, How & Why of IPv6 Transit and Peering

- Relationships
  - Peers, Transit providers, agreements
- Technology
  - BGP peering
  - BGP filtering

## IPv6 Transit and Peering Relationships

- Who should I Peer with?
  - Everyone you Peer with on IPv4
- Who should I buy Transit from?
  - Whoever you buy IPv4 Transit from
- What terms should I negotiate?
  - The same terms as are in your IPv4 agreements
- Revisit this once IPv6 traffic levels come up

## Carpe Diem?

- This may be an opportunity to expand your peering
- Acquire new IPv6-only peers
  - In addition to IPv4-only and dual-stack peers
- The faster you build out IPv6, the more likely this is

IPv6 Peering and Transit

# TECHNOLOGY

## IPv6 BGP Peering

- Assuming existing IPv4 BGP sessions
  - Not a greenfield
  - eBGP and iBGP already established



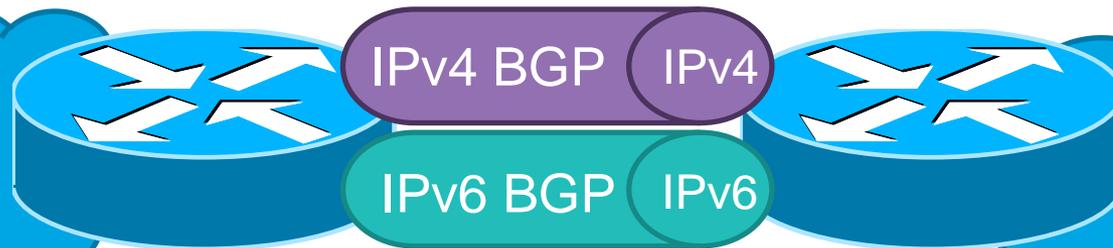
## Single Session

- Exchange IPv6 routes over existing IPv4 peering?
  - Single BGP session between peers
  - Leverages Multiprotocol BGP (MP-BGP / RFC 4760)
    - *Adding IPv6 NLRI bounces the session*
  - IPv4 and IPv6 have a shared fate



## Dual Session

- Establish new, IPv6-only peering?
  - Two BGP sessions between peers
  - IPv6 topology independent of IPv4, and vice versa
    - Outages: IPv6 session goes down if IPv6 reachability is lost
    - Maintenance: IPv4 and IPv6 sessions don't affect each other
    - Operational Clarity



## Junos Configuration Example

```
routing-options {  
    router-id 10.10.10.5  
    autonomous-system 65000;  
}  
protocols bgp group v6_PEERS {  
    type external;  
    neighbor 2001:db8:8000:4200::1 {  
        peer-as 64109;  
    }  
}
```

## IOS Configuration Example

```
router bgp 65000  
  bgp router-id 10.10.10.5  
  no bgp default ipv4-unicast
```

***NOTE: When you add “address-family” the BGP config splits into three parts: ipv4, ipv6, and the “main” BGP config.***

!

```
  address-family ipv6  
    neighbor 2001:DB8:8000:4200::1 activate  
    no synchronization  
  exit-address-family
```

# IPv6 BGP Session Verification Commands

Juniper Command	Cisco IOS Command	Co-Ordinating Definition
show ipv6 neighbors	show ipv6 neighbors	Show IPv6 neighbor cache information
	show ipv6 protocols	Show IPv6 Routing Protocols
show interface	show ipv6 interface	Show IPv6 interface information
show route protocol bgp table inet6.0	show bgp all show bgp ipv6 unicast	Show IPv6 unicast routes [Network, Next Hop, Metric, LocPrf, Weight, AS Path, etc.]
show bgp summary	show bgp ipv6 unicast summary	Show IPv6 unicast peers [AS, Up/Down, # of learned prefixes, etc.]
show bgp neighbor	show bgp ipv6 unicast neighbors	Show detailed information about each peer
show route advertising-protocol bgp <ipv6 addr>	show bgp ipv6 unicast neighbors [<ipv6 addr> advertised-routes]	Show IPv6 prefixes advertised to a peer
show route receive-protocol bgp <ipv6 addr> table inet6.0	show bgp ipv6 unicast neighbors [<ipv6 addr> recieved-routes]	Show IPv6 prefixes received from a peer – requires “neighbor soft-reconfiguration”
show route <prefix>/<length>	show bgp ipv6 <prefix>/<length>	Show information about a specific IPv6 prefix
show policy <policy-name>	show ipv6 prefix-list [summary detail] show ipv6 prefix-list <name>	Show IPv6 prefix-list information
Use traceoptions with flags	debug bgp ipv6 unicast	Debug BGP IPv6 packets
clear bgp neighbor <neighbor>	clear bgp ipv6 unicast *	Clear the BGP Session to all peers
ping [ipv6]	ping [ipv6]	Send echo messages
traceroute [ipv6]	traceroute [ipv6]	Trace route to destination

## IPv6 BGP Filtering

- Methodology not special, generally mirrors IPv4
  - Explicit, bogons/martians, maximum-prefix, prefix size, etc.
  - The devil's in the details
- Three high-level “themes:”
  - Filtering customers
  - Filtering Peers and Transit providers
  - Filtering your own routes

## Filtering Routes Coming From Customers

- Explicit filter:
  - Allow customer network(s)
  - Deny all else
  - Customer is responsible for updating you
    - Should be infrequent if at all in IPv6
- Prefix size?
  - In IPv4 it's common to allow down to /32 (for blackholing, etc.)
  - Orders of magnitude more addresses in v6
  - If allowed in IPv6 (/128); **ensure proper maximum-prefix limits**
  - Filter out more specifics before announcing to Peer / Transit
    - Could filter to /48 if customer needs to leverage more specifics

## Ingress Customer Prefix Filter Example

- Assume your customer's prefix is 2001:db8::/32
- IOS:

```
ipv6 prefix-list ipv6-from-customer permit 2001:db8::/32  
ipv6 prefix-list ipv6-from-customer deny 0::/0 le 128
```

- Junos:

```
policy-statement ipv6-from-customer {  
  from {  
    family inet6;  
    route-filter 2001:db8::/32 exact next policy;  
  }  
  then reject;  
}
```

## Filtering Routes Coming From Peers & Upstreams

### 1. Generate prefix filters from IRR:

- Everyone should register in an Internet Routing Registry (IRR)
- All
- As
- 

***NOTE: The number of IPv6 prefixes being announced from Peers and especially from Upstreams is growing rapidly. Be prepared to monitor and update your limits frequently.***

### 2. Bogon and max-prefix

- Some networks don't use an IRR (yell at them)
- Bogon filter rejects the crazy stuff (reserved, small prefix, etc.)
- A maximum-prefix limit protects against route overload
- Better than nothing

## IPv6 Bogon / Martian Filtering Basics

Prefix	Description	Action
::/0	Default Route	Deny, Exact
::/8	Contains Loop back Address (::1/128), Unspecified Address (::/128), IETF reserved Address (formerly IPv4-compatible IPv6 Address) (::/96), and IPv4-mapped IPv6 Address (::ffff:0:0/96)	Deny, Or Longer
2000::/3	Global Unicast	Allow, /48 Or Shorter
2001::/32	Toronto	Allow, Exact
2001:db8::/32	Documentation Address	Deny, Or Longer
2002::/16	6to4	Allow, Exact
3ffe::/16	Former 6bone	Deny, Or Longer
fe80::/10	Link-local Unicast	Deny, Or Longer
fec0::/10	IETF reserved Address (formerly Site-local Address)	Deny, Or Longer
fc00::/7	Unique-local Address	Deny, Or Longer
ff00::/8	Multicast Address	Deny, Or Longer

## Example “Loose” Bogon Filter - IOS

```
ipv6 prefix-list ipv6-bogon-loose deny 0::/0
ipv6 prefix-list ipv6-bogon-loose deny 0::/8 le 128
ipv6 prefix-list ipv6-bogon-loose permit 2000::/3 le 48
ipv6 prefix-list ipv6-bogon-loose deny <your prefix> le 128
ipv6 prefix-list ipv6-bogon-loose permit 2001::/32
ipv6 prefix-list ipv6-bogon-loose deny 2001::/32 le 128
ipv6 prefix-list ipv6-bogon-loose deny 2001:db8::/32 le 128
ipv6 prefix-list ipv6-bogon-loose permit 2002::/16
ipv6 prefix-list ipv6-bogon-loose deny 2002::/16 le 128
ipv6 prefix-list ipv6-bogon-loose deny 3ffe::/16 le 128
ipv6 prefix-list ipv6-bogon-loose deny fe80::/10 le 128
ipv6 prefix-list ipv6-bogon-loose deny fec0::/10 le 128
ipv6 prefix-list ipv6-bogon-loose deny fc00::/7 le 128
ipv6 prefix-list ipv6-bogon-loose deny ff00::/8 le 128
ipv6 prefix-list ipv6-bogon-loose deny 0::/0 le 128
```

## Example “Loose” Bogon Filter - Junos

```
policy-statement ipv6-bogon-loose {  
  from {  
    family inet6;  
    route-filter 0::/0 exact;  
    route-filter 0::/8 orlonger;  
    route-filter 2000::/3 prefix-length-range /3-/48 next policy;  
    route-filter <your prefix> orlonger next policy;  
    route-filter 2001::/32 exact next policy;  
    route-filter 2001::/32 longer;  
    route-filter 2001:db8::/32 orlonger;  
    route-filter 2002::/16 exact next policy;  
    route-filter 2002::/16 longer;  
    route-filter 3ffe::/16 orlonger;  
    route-filter fe80::/10 orlonger;  
    route-filter fec0::/10 orlonger;  
    route-filter fc00::/7 orlonger;  
    route-filter ff00::/8 orlonger;  
    route-filter 0::/0 orlonger;  
  }  
  then {  
    reject;  
  }  
}
```

## Strict Bogon Filters

- Not all of 2000::/3 has been allocated to RIRs
  - <http://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xml>
- Other prefixes may be considered “bogon”
  - <http://www.team-cymru.org/Services/Bogons/fullbogons-ipv6.txt>
- RIRs hand out different space in different sized blocks
  - Provider Assigned (PA) vs Provider Independent (PI) space
  - Micro-allocations, etc.
- Strict filters must be actively maintained
  - Constantly changing landscape by design

## Filtering Your Own Routes

- Only allow you and your customer's space
- Don't deaggregate unnecessarily
  - Only /48 and shorter (larger) prefixes
- Use communities to identify and filter routes
  - Level of hierarchy – Filter more specifics
  - Type of customer – BGP customers more likely to multihome
  - Use – No need to advertise specific infrastructure prefixes
  - Locational – Allows regionalization for TE, etc.

## Some Parting Thoughts

- Take the opportunity to get it right!
- Use an IRR
  - Register in an existing database
  - Register your own database
  - Synchronize your own filters with your IRR
- Consider using a comprehensive IPAM solution
  - Massive address space
  - DNS/DHCP complexities
- Consider registering in the PeeringDB

## A Few References

- BGP Configuration:
  - <http://www.scribd.com/doc/54134779/IPv6-Routing-Grundemann-Hughes-Sexton>
- Route Filtering:
  - <http://www.team-cymru.org/ReadingRoom/Templates/IPv6Routers/>
  - <http://www.space.net/~gert/RIPE/ipv6-filters.html>
- IRR:
  - <http://www.irr.net/>
  - <https://www.isc.org/software/irrtoolset>

# Questions?

[@ChrisGrundemann](#)  
[c.grundemann@cablelabs.com](mailto:c.grundemann@cablelabs.com)  
<http://gogonet.gogo6.com/profile/cgrundemann>